

09/954,719


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide

THE ACM DIGITAL LIBRARY
[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used method for selecting web pages dimensional model generating query result match query value fact table produce join result

Found 25 of 198,617

 Sort results
by

 relevance
 Save results to a Binder

[Try an Advanced Search](#)

 Display
results

 expanded form
 Search Tips

[Try this search in The ACM Guide](#)
 Open results in a new window

Results 1 - 20 of 25

Result page: 1 [2](#) [next](#)
 Relevance scale

- 1 [Data warehouse models: Dimensional modeling: identifying, classifying & applying patterns](#)

Mary Elizabeth Jones, Il-Yeol Song
November 2005 **Proceedings of the 8th ACM international workshop on Data warehousing and OLAP DOLAP '05**

Publisher: ACM Press

Full text available: [pdf\(289.04 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Software design is a complex activity. A successful designer requires knowledge and training in specific design techniques combined with practical experience. Designing a dimensional model embodies this challenge. This paper presents Dimensional Design Patterns (DDPs) and their applications to the design of dimensional models. We describe a metamodel of the DDPs and show their integration into Kimball's dimensional modeling design process so they can be applied to design problems using a known p ...

Keywords: data warehouse, dimensional modeling, patterns, software engineering

- 2 [Book reviews: Review of The data warehouse toolkit: the complete guide to dimensional modeling \(2nd edition\) by Ralph Kimball, Margy Ross. John Wiley & Sons, Inc. 2002.](#)

Alexander A. Anisimov
September 2003 **ACM SIGMOD Record**, Volume 32 Issue 3

Publisher: ACM Press

Full text available: [pdf\(22.79 KB\)](#) Additional Information: [full citation](#)

- 3 [Modeling strategies and alternatives for data warehousing projects](#)

Nenad Jukic
April 2006 **Communications of the ACM**, Volume 49 Issue 4

Publisher: ACM Press

Full text available: [pdf\(156.15 KB\)](#) [html\(27.55 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Choosing the appropriate modeling approach is often the critical factor in the success or

failure of a data warehousing implementation.

4 Papers: Dimensional modeling for a data warehouse

 Ashish Kumar Dash, Rakesh Agarwal
November 2001 **ACM SIGSOFT Software Engineering Notes**, Volume 26 Issue 6

Publisher: ACM Press

Full text available:  pdf(231.45 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

A multidimensional database stores data as groups of field category values into dimensions, and then groups these dimensions into multidimensional arrays. Specific field category values that may occur in data identify either the rows or columns of array dimensions. The specific grouped field categories themselves identify the row or column array dimensions. This view, when presented to the end user, bring in more relevance and business sense for practical decision making than the views presented ...

5 Why decision support fails and how to fix it

 Ralph Kimball, Kevin Strehlo
September 1995 **ACM SIGMOD Record**, Volume 24 Issue 3

Publisher: ACM Press

Full text available:  pdf(722.24 KB) Additional Information: [full citation](#), [citations](#), [index terms](#)

6 A comparison of data warehousing methodologies

 Arun Sen, Atish P. Sinha
March 2005 **Communications of the ACM**, Volume 48 Issue 3

Publisher: ACM Press

Full text available:  pdf(117.81 KB)  html(28.41 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Using a common set of attributes to determine which methodology to use in a particular data warehousing project.

7 Capturing summarizability with integrity constraints in OLAP

 Carlos A. Hurtado, Claudio Gutierrez, Alberto O. Mendelzon
September 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 3

Publisher: ACM Press

Full text available:  pdf(710.79 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

In multidimensional data models intended for online analytic processing (OLAP), data are viewed as points in a multidimensional space. Each dimension has structure, described by a directed graph of categories, a set of members for each category, and a child/parent relation between members. An important application of this structure is to use it to infer summarizability, that is, whether an aggregate view defined for some category can be correctly derived from a set of precomputed views defined f ...

Keywords: OLAP, data warehousing, integrity constraints, query-optimization, summarizability

8 Query and view processing: Aggregate queries in peer-to-peer OLAP

 Mauricio Minuto Espil, Alejandro A. Vaisman
November 2004 **Proceedings of the 7th ACM international workshop on Data warehousing and OLAP DOLAP '04**

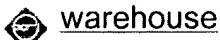
Publisher: ACM Press

Full text available: [pdf\(375.80 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A peer-to-peer (P2P) data management system consists essentially in a network of peer systems, each maintaining full autonomy over its own data resources. Data exchange between peers occurs when one of them, in the role of a local peer, needs data available in other nodes, denoted the acquaintances of the local peer. No global schema is assumed to exist for any data under this computing paradigm. Henceforth, data provided by an acquaintance of a local peer must be adapted, in a manner that an ...

Keywords: OLAP, P2P computing, data warehousing

9 Aggregation everywhere: data reduction and transformation in the Phoenix data



warehouse

Steven Tolkin

November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP DOLAP '99**

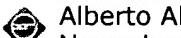
Publisher: ACM Press

Full text available: [pdf\(1.23 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the Phoenix system, which loads a data warehouse and then reports against it. Between the raw atomic data of the source system and the business measures presented to users there are many computing environments. Aggregation occurs everywhere: initial bucketing by the natural keys on the mainframe, loading the fact table using a mapping table, maintaining aggregate tables and reporting tables in the data base, in the GUI, in SQL queries issued on behalf of client tools by ...

Keywords: OLAP, SQL, aggregation, data lineage, data warehouse

10 Query processing: Implementing operations to navigate semantic star schemas



Alberto Abelló, José Samos, Fèlix Saltor

November 2003 **Proceedings of the 6th ACM international workshop on Data warehousing and OLAP DOLAP '03**

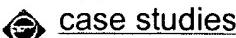
Publisher: ACM Press

Full text available: [pdf\(193.82 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In the last years, lots of work have been devoted to multidimensional modeling, star shape schemas and OLAP operations. However, "drill-across" has not captured as much attention as other operations. This operation allows to change the subject of analysis keeping the same analysis space we were using to analyze another subject. It is assumed that this can be done if both subjects share exactly the same analysis dimensions. In this paper, besides the implementation of an algebraic set of operatio ...

Keywords: OLAP operations, SQL, drill-across, semantic relationships, star schema

11 Approaches to the development of multi-dimensional databases: lessons from four



case studies

Helen Hasan, Peter Hyland, David Dodds, Raja Veeraraghavan

May 2000 **ACM SIGMOD Database**, Volume 31 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.21 MB\)](#) Additional Information: [full citation](#), [index terms](#)

Keywords: OLAP, case study, executive information systems, management, methodology, multi-multidimensional database, organization, relational database

12 Approaches to the development of multi-dimensional databases: lessons from four

case studies

Helen Hasan, Peter Hyland, David Dodds, Raja Veeraraghavan

June 2000 **ACM SIGMIS Database**, Volume 31 Issue 3

Publisher: ACM Press

Full text available: [pdf\(1.21 MB\)](#) Additional Information: [full citation](#), [abstract](#), [index terms](#)

The paper explores the manner in which an organization's data and information can be effectively utilized to assist an organization to achieve its business objectives. With the increased popularity of data warehousing and executive information systems, there is renewed interest by IT practitioners in data models and database structures, in particular multi-dimensional forms, which have joined their relational counterparts as legitimate tools for extracting vital business information from an orga ...

Keywords: OLAP, case study, executive information systems, management, methodology, multi-dimensional database, organization, relational database

13 Deriving initial data warehouse structures from the conceptual data models of the

underlying operational information systems

Michael Boehnlein, Achim Ulbrich-vom Ende

November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP DOLAP '99**

Publisher: ACM Press

Full text available: [pdf\(1.40 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years the construction of large scale data schemes for operational systems has been the major problem of conceptual data modeling for business needs. Multidimensional data structures used for decision support applications in data warehouses have rather different requirements to data modeling techniques. In case of operational systems the data models are created from application specific requirements. The data models in data warehouses base on the analytical requirements of the use ...

Keywords: conceptual data model, data warehouse, decision support system, entity relationship model (ERM), snowflake scheme, star scheme, structured entity relationship model (SERM)

14 Conceptual modeling for ETL processes

Panos Vassiliadis, Alkis Simitsis, Spiros Skiadopoulos

November 2002 **Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP DOLAP '02**

Publisher: ACM Press

Full text available: [pdf\(471.56 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Extraction-Transformation-Loading (ETL) tools are pieces of software responsible for the extraction of data from several sources, their cleansing, customization and insertion into a data warehouse. In this paper, we focus on the problem of the definition of ETL activities and provide formal foundations for their conceptual representation. The proposed conceptual model is (a) customized for the tracing of inter-attribute relationships and the respective ETL activities in the early stages of a dat ...

Keywords: ETL, conceptual modeling, data warehousing

15 Data Warehouse: Index filtering and view materialization in ROLAP environment

Shi Guang Qiu, Tok Wang Ling

October 2001 **Proceedings of the tenth international conference on Information and knowledge management CIKM '01**

Publisher: ACM Press

Full text available:  pdf(1.17 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Using materialized view to accelerate OLAP queries is one of the most common methods used in ROLAP systems. However, high storage and computation cost make this method very difficult to be implemented in the actual environment. Among various issues associated with this, index selection and view materialization are two of the top challenges. In this paper, we propose to build indexes on subsets of the primary keys rather than the full sets if the index selectivity for these smaller indexes can be ...

16 An alternative storage organization for ROLAP aggregate views based on cubetrees

Yannis Kotidis, Nick Roussopoulos

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98**, Volume 27 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.19 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Relational On-Line Analytical Processing (ROLAP) is emerging as the dominant approach in data warehousing with decision support applications. In order to enhance query performance, the ROLAP approach relies on selecting and materializing in summary tables appropriate subsets of aggregate views which are then engaged in speeding up OLAP queries. However, a straight forward relational storage implementation of materialized ROLAP views is immensely wasteful on storage and incredibly inadeq ...

17 Research session 6: OLAP and constraints: OLAP dimension constraints

Carlos A. Hurtado, Alberto O. Mendelzon

June 2002 **Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems PODS '02**

Publisher: ACM Press

Full text available:  pdf(186.06 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In multidimensional data models intended for online analytic processing (OLAP), data are viewed as points in a multidimensional space. Each dimension has structure, described by a directed graph of categories, a set of members for each category, and a child/parent relation between members. An important application of this structure is to use it to infer summarizability, that is, whether an aggregate view defined for some category can be correctly derived from a set of precomputed views defined f ...

18 Data warehouse construction: Building a web warehouse for accessibility data

Christian Thomsen, Torben Bach Pedersen

November 2006 **Proceedings of the 9th ACM international workshop on Data warehousing and OLAP DOLAP '06**

Publisher: ACM Press

Full text available:  pdf(181.96 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As more and more information is available on the web, it is a problem that many web resources are not accessible, i.e., are not usable for users with special needs. For

example, for a web page to be accessible, it should give text alternatives (i.e., explanatory texts) for images such that blind users that have the web pages read aloud automatically also can obtain information about the images. In the European Internet Accessibility Observatory (EIAO) project, a crawler that will evaluate the ac ...

Keywords: accessibility, web warehouse

19 [Reports: Report on the ACM fourth international workshop on data warehousing and](#) 

 [OLAP \(DOLAP 2001\)](#)

Joachim Hammer

June 2002 **ACM SIGMOD Record**, Volume 31 Issue 2

Publisher: ACM Press

Full text available:  pdf(413.34 KB) Additional Information: [full citation](#), [abstract](#)

The Fourth Annual ACM International Workshop on Data Warehousing and Online Analytical Processing (DOLAP 2001) was held in Atlanta, GA, USA, in November 2001, in conjunction with the Tenth International Conference on Information and Knowledge Management (CIKM 2001). Although this was only the fourth annual meeting, DOLAP has already become an important and broadly accepted forum for researchers and practitioners to share their findings in theoretical foundations, current methodologies, practical ...

20 [A method for developing dimensional data marts](#) 

 Tim Chenoweth, David Schuff, Robert St. Louis

December 2003 **Communications of the ACM**, Volume 46 Issue 12

Publisher: ACM Press

Full text available:  pdf(104.91 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index](#)
 html(26.61 KB) [terms](#)

Decision-oriented dimensional data marts are fundamentally different than transaction-oriented relational databases. A distinctive methodology and a different set of tools are required for their effective development.

Results 1 - 20 of 25

Result page: [1](#) [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

091954,719


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [system for selecting web pages](#) [dimensional model](#) [generating query result](#) [match query](#) [value fact table](#) [produce join result](#)

Found 25 of 198,617

Sort results
by

[Save results to a Binder](#)
[Try an Advanced Search](#)

Display
results

[Search Tips](#)
[Try this search in The ACM Guide](#)
[Open results in a new window](#)

Results 1 - 20 of 25

Result page: [1](#) [2](#) [next](#)

Relevance scale

- 1** [Data warehouse models: Dimensional modeling: identifying, classifying & applying patterns](#)

Mary Elizabeth Jones, Il-Yeol Song
November 2005 **Proceedings of the 8th ACM international workshop on Data warehousing and OLAP DOLAP '05**

Publisher: ACM Press

Full text available: [pdf\(289.04 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Software design is a complex activity. A successful designer requires knowledge and training in specific design techniques combined with practical experience. Designing a dimensional model embodies this challenge. This paper presents Dimensional Design Patterns (DDPs) and their applications to the design of dimensional models. We describe a metamodel of the DDPs and show their integration into Kimball's dimensional modeling design process so they can be applied to design problems using a known p ...

Keywords: data warehouse, dimensional modeling, patterns, software engineering

- 2** [Book reviews: Review of The data warehouse toolkit: the complete guide to dimensional modeling \(2nd edition\) by Ralph Kimball, Margy Ross. John Wiley & Sons, Inc. 2002.](#)

Alexander A. Anisimov
September 2003 **ACM SIGMOD Record**, Volume 32 Issue 3

Publisher: ACM Press

Full text available: [pdf\(22.79 KB\)](#) Additional Information: [full citation](#)

- 3** [Modeling strategies and alternatives for data warehousing projects](#)

Nenad Jukic
April 2006 **Communications of the ACM**, Volume 49 Issue 4

Publisher: ACM Press

Full text available: [pdf\(156.15 KB\)](#) [html\(27.55 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Choosing the appropriate modeling approach is often the critical factor in the success or

failure of a data warehousing implementation.

4 Papers: Dimensional modeling for a data warehouse

 Ashish Kumar Dash, Rakesh Agarwal
November 2001 **ACM SIGSOFT Software Engineering Notes**, Volume 26 Issue 6

Publisher: ACM Press

Full text available:  pdf(231.45 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

A multidimensional database stores data as groups of field category values into dimensions, and then groups these dimensions into multidimensional arrays. Specific field category values that may occur in data identify either the rows or columns of array dimensions. The specific grouped field categories themselves identify the row or column array dimensions. This view, when presented to the end user, bring in more relevance and business sense for practical decision making than the views presented ...

5 Why decision support fails and how to fix it

 Ralph Kimball, Kevin Strehlo
September 1995 **ACM SIGMOD Record**, Volume 24 Issue 3

Publisher: ACM Press

Full text available:  pdf(722.24 KB) Additional Information: [full citation](#), [citations](#), [index terms](#)

6 A comparison of data warehousing methodologies

 Arun Sen, Atish P. Sinha
March 2005 **Communications of the ACM**, Volume 48 Issue 3

Publisher: ACM Press

Full text available:  pdf(117.81 KB)  html(28.41 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Using a common set of attributes to determine which methodology to use in a particular data warehousing project.

7 Capturing summarizability with integrity constraints in OLAP

 Carlos A. Hurtado, Claudio Gutierrez, Alberto O. Mendelzon
September 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 3

Publisher: ACM Press

Full text available:  pdf(710.79 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

In multidimensional data models intended for online analytic processing (OLAP), data are viewed as points in a multidimensional space. Each dimension has structure, described by a directed graph of categories, a set of members for each category, and a child/parent relation between members. An important application of this structure is to use it to infer summarizability, that is, whether an aggregate view defined for some category can be correctly derived from a set of precomputed views defined f ...

Keywords: OLAP, data warehousing, integrity constraints, query-optimization, summarizability

8 Query and view processing: Aggregate queries in peer-to-peer OLAP

 Mauricio Minuto Espil, Alejandro A. Vaisman
November 2004 **Proceedings of the 7th ACM international workshop on Data warehousing and OLAP DOLAP '04**

Publisher: ACM Press

Full text available: [A.pdf\(375.80 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A peer-to-peer (P2P) data management system consists essentially in a network of peer systems, each maintaining full autonomy over its own data resources. Data exchange between peers occurs when one of them, in the role of a local peer, needs data available in other nodes, denoted the acquaintances of the local peer. No global schema is assumed to exist for any data under this computing paradigm. Henceforth, data provided by an acquaintance of a local peer must be adapted, in a manner that an ...

Keywords: OLAP, P2P computing, data warehousing

9 Aggregation everywhere: data reduction and transformation in the Phoenix data



Steven Tolkin

November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP DOLAP '99**

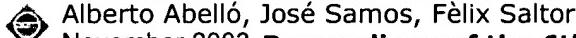
Publisher: ACM Press

Full text available: [A.pdf\(1.23 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the Phoenix system, which loads a data warehouse and then reports against it. Between the raw atomic data of the source system and the business measures presented to users there are many computing environments. Aggregation occurs everywhere: initial bucketing by the natural keys on the mainframe, loading the fact table using a mapping table, maintaining aggregate tables and reporting tables in the data base, in the GUI, in SQL queries issued on behalf of client tools by ...

Keywords: OLAP, SQL, aggregation, data lineage, data warehouse

10 Query processing: Implementing operations to navigate semantic star schemas



Alberto Abelló, José Samos, Fèlix Saltor

November 2003 **Proceedings of the 6th ACM international workshop on Data warehousing and OLAP DOLAP '03**

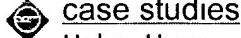
Publisher: ACM Press

Full text available: [A.pdf\(193.82 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In the last years, lots of work have been devoted to multidimensional modeling, star shape schemas and OLAP operations. However, "drill-across" has not captured as much attention as other operations. This operation allows to change the subject of analysis keeping the same analysis space we were using to analyze another subject. It is assumed that this can be done if both subjects share exactly the same analysis dimensions. In this paper, besides the implementation of an algebraic set of operatio ...

Keywords: OLAP operations, SQL, drill-across, semantic relationships, star schema

11 Approaches to the development of multi-dimensional databases: lessons from four



Helen Hasan, Peter Hyland, David Dodds, Raja Veeraraghavan

May 2000 **ACM SIGMIS Database**, Volume 31 Issue 2

Publisher: ACM Press

Full text available: [A.pdf\(1.21 MB\)](#) Additional Information: [full citation](#), [index terms](#)

Keywords: OLAP, case study, executive information systems, management, methodology, multi-multidimensional database, organization, relational database

12 Approaches to the development of multi-dimensional databases: lessons from four case studies

 Helen Hasan, Peter Hyland, David Dodds, Raja Veeraraghavan
June 2000 **ACM SIGMIS Database**, Volume 31 Issue 3

Publisher: ACM Press

Full text available:  pdf(1.21 MB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

The paper explores the manner in which an organization's data and information can be effectively utilized to assist an organization to achieve its business objectives. With the increased popularity of data warehousing and executive information systems, there is renewed interest by IT practitioners in data models and database structures, in particular multi-dimensional forms, which have joined their relational counterparts as legitimate tools for extracting vital business information from an orga ...

Keywords: OLAP, case study, executive information systems, management, methodology, multi-dimensional database, organization, relational database

13 Deriving initial data warehouse structures from the conceptual data models of the underlying operational information systems

 Michael Boehnlein, Achim Ulbrich-vom Ende
November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP DOLAP '99**

Publisher: ACM Press

Full text available:  pdf(1.40 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years the construction of large scale data schemes for operational systems has been the major problem of conceptual data modeling for business needs. Multidimensional data structures used for decision support applications in data warehouses have rather different requirements to data modeling techniques. In case of operational systems the data models are created from application specific requirements. The data models in data warehouses base on the analytical requirements of the use ...

Keywords: conceptual data model, data warehouse, decision support system, entity relationship model (ERM), snowflake scheme, star scheme, structured entity relationship model (SERM)

14 Conceptual modeling for ETL processes

 Panos Vassiliadis, Alkis Simitsis, Spiros Skiadopoulos
November 2002 **Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP DOLAP '02**

Publisher: ACM Press

Full text available:  pdf(471.56 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Extraction-Transformation-Loading (ETL) tools are pieces of software responsible for the extraction of data from several sources, their cleansing, customization and insertion into a data warehouse. In this paper, we focus on the problem of the definition of ETL activities and provide formal foundations for their conceptual representation. The proposed conceptual model is (a) customized for the tracing of inter-attribute relationships and the respective ETL activities in the early stages of a dat ...

Keywords: ETL, conceptual modeling, data warehousing

15 Data Warehouse: Index filtering and view materialization in ROLAP environment

 Shi Guang Qiu, Tok Wang Ling

October 2001 **Proceedings of the tenth international conference on Information and knowledge management CIKM '01**

Publisher: ACM Press

Full text available:  pdf(1.17 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Using materialized view to accelerate OLAP queries is one of the most common methods used in ROLAP systems. However, high storage and computation cost make this method very difficult to be implemented in the actual environment. Among various issues associated with this, index selection and view materialization are two of the top challenges. In this paper, we propose to build indexes on subsets of the primary keys rather than the full sets if the index selectivity for these smaller indexes can be ...

16 An alternative storage organization for ROLAP aggregate views based on cubetrees

 Yannis Kotidis, Nick Roussopoulos

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98**, Volume 27 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.19 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Relational On-Line Analytical Processing (ROLAP) is emerging as the dominant approach in data warehousing with decision support applications. In order to enhance query performance, the ROLAP approach relies on selecting and materializing in summary tables appropriate subsets of aggregate views which are then engaged in speeding up OLAP queries. However, a straight forward relational storage implementation of materialized ROLAP views is immensely wasteful on storage and incredibly inadequate ...

17 Research session 6: OLAP and constraints: OLAP dimension constraints

 Carlos A. Hurtado, Alberto O. Mendelzon

June 2002 **Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems PODS '02**

Publisher: ACM Press

Full text available:  pdf(186.06 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In multidimensional data models intended for online analytic processing (OLAP), data are viewed as points in a multidimensional space. Each dimension has structure, described by a directed graph of categories, a set of members for each category, and a child/parent relation between members. An important application of this structure is to use it to infer summarizability, that is, whether an aggregate view defined for some category can be correctly derived from a set of precomputed views defined for ...

18 Data warehouse construction: Building a web warehouse for accessibility data

 Christian Thomsen, Torben Bach Pedersen

November 2006 **Proceedings of the 9th ACM international workshop on Data warehousing and OLAP-DOLAP '06**

Publisher: ACM Press

Full text available:  pdf(181.96 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As more and more information is available on the web, it is a problem that many web resources are not accessible, i.e., are not usable for users with special needs. For

example, for a web page to be accessible, it should give text alternatives (i.e., explanatory texts) for images such that blind users that have the web pages read aloud automatically also can obtain information about the images. In the European Internet Accessibility Observatory (EIAO) project, a crawler that will evaluate the ac ...

Keywords: accessibility, web warehouse

19 Reports: Report on the ACM fourth international workshop on data warehousing and 

 OLAP (DOLAP 2001)

Joachim Hammer

June 2002 **ACM SIGMOD Record**, Volume 31 Issue 2

Publisher: ACM Press

Full text available:  pdf(413.34 KB) Additional Information: [full citation](#), [abstract](#)

The Fourth Annual ACM International Workshop on Data Warehousing and Online Analytical Processing (DOLAP 2001) was held in Atlanta, GA, USA, in November 2001, in conjunction with the Tenth International Conference on Information and Knowledge Management (CIKM 2001). Although this was only the fourth annual meeting, DOLAP has already become an important and broadly accepted forum for researchers and practitioners to share their findings in theoretical foundations, current methodologies, practical ...

20 A method for developing dimensional data marts 

 Tim Chenoweth, David Schuff, Robert St. Louis

December 2003 **Communications of the ACM**, Volume 46 Issue 12

Publisher: ACM Press

Full text available:  pdf(104.91 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index](#)
 html(26.61 KB) [terms](#)

Decision-oriented dimensional data marts are fundamentally different than transaction-oriented relational databases. A distinctive methodology and a different set of tools are required for their effective development.

Results 1 - 20 of 25

Result page: [1](#) [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

09/05 4,719


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide

 "selecting web" + "dimension model" + "generating query result"

THE ACM DIGITAL LIBRARY
[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used **selecting web dimension model generating query result match query value fact table produce join result**

Found 25 of 198,617

Sort results by

[Save results to a Binder](#)

Display results

[Search Tips](#)
 Open results in a new window

 Try an [Advanced Search](#)
Try this search in [The ACM Guide](#)

Results 1 - 20 of 25

Result page: **1** [2](#) [next](#)

Relevance scale



- 1** [Data warehouse models: Dimensional modeling: identifying, classifying & applying patterns](#)



Mary Elizabeth Jones, Il-Yeol Song

November 2005 **Proceedings of the 8th ACM international workshop on Data warehousing and OLAP DOLAP '05****Publisher:** ACM PressFull text available: [pdf\(289.04 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Software design is a complex activity. A successful designer requires knowledge and training in specific design techniques combined with practical experience. Designing a dimensional model embodies this challenge. This paper presents Dimensional Design Patterns (DDPs) and their applications to the design of dimensional models. We describe a metamodel of the DDPs and show their integration into Kimball's dimensional modeling design process so they can be applied to design problems using a known p ...

Keywords: data warehouse, dimensional modeling, patterns, software engineering



- 2** [Book reviews: Review of The data warehouse toolkit: the complete guide to dimensional modeling \(2nd edition\) by Ralph Kimball, Margy Ross. John Wiley & Sons, Inc. 2002.](#)



Alexander A. Anisimov

September 2003 **ACM SIGMOD Record**, Volume 32 Issue 3**Publisher:** ACM PressFull text available: [pdf\(22.79 KB\)](#) Additional Information: [full citation](#)

- 3** [Modeling strategies and alternatives for data warehousing projects](#)



Nenad Jukic

April 2006 **Communications of the ACM**, Volume 49 Issue 4**Publisher:** ACM PressFull text available: [pdf\(156.15 KB\)](#) [html\(27.55 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Choosing the appropriate modeling approach is often the critical factor in the success or failure of a data warehousing implementation.

4 Papers: Dimensional modeling for a data warehouse □

 Ashish Kumar Dash, Rakesh Agarwal
November 2001 **ACM SIGSOFT Software Engineering Notes**, Volume 26 Issue 6

Publisher: ACM Press

Full text available: [pdf\(231.45 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

A multidimensional database stores data as groups of field category values into dimensions, and then groups these dimensions into multidimensional arrays. Specific field category values that may occur in data identify either the rows or columns of array dimensions. The specific grouped field categories themselves identify the row or column array dimensions. This view, when presented to the end user, bring in more relevance and business sense for practical decision making than the views presented ...

5 Why decision support fails and how to fix it □

 Ralph Kimball, Kevin Strehlo
September 1995 **ACM SIGMOD Record**, Volume 24 Issue 3

Publisher: ACM Press

Full text available: [pdf\(722.24 KB\)](#) Additional Information: [full citation](#), [citations](#), [index terms](#)

6 A comparison of data warehousing methodologies □

 Arun Sen, Atish P. Sinha
March 2005 **Communications of the ACM**, Volume 48 Issue 3

Publisher: ACM Press

Full text available: [pdf\(117.81 KB\)](#) [html\(28.41 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Using a common set of attributes to determine which methodology to use in a particular data warehousing project.

7 Capturing summarizability with integrity constraints in OLAP □

 Carlos A. Hurtado, Claudio Gutierrez, Alberto O. Mendelzon
September 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 3

Publisher: ACM Press

Full text available: [pdf\(710.79 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#), [review](#)

In multidimensional data models intended for online analytic processing (OLAP), data are viewed as points in a multidimensional space. Each dimension has structure, described by a directed graph of categories, a set of members for each category, and a child/parent relation between members. An important application of this structure is to use it to infer summarizability, that is, whether an aggregate view defined for some category can be correctly derived from a set of precomputed views defined f ...

Keywords: OLAP, data warehousing, integrity constraints, query-optimization, summarizability

8 Query and view processing: Aggregate queries in peer-to-peer OLAP □

 Mauricio Minuto Espil, Alejandro A. Vaisman
November 2004 **Proceedings of the 7th ACM international workshop on Data warehousing and OLAP DOLAP '04**

Publisher: ACM Press

Full text available: Additional Information:

 pdf(375.80 KB)

[full citation](#), [abstract](#), [references](#), [index terms](#)

A peer-to-peer (P2P) data management system consists essentially in a network of peer systems, each maintaining full autonomy over its own data resources. Data exchange between peers occurs when one of them, in the role of a local peer, needs data available in other nodes, denoted the acquaintances of the local peer. No global schema is assumed to exist for any data under this computing paradigm. Henceforth, data provided by an acquaintance of a local peer must be adapted, in a manner that an ...

Keywords: OLAP, P2P computing, data warehousing

9 Aggregation everywhere: data reduction and transformation in the Phoenix data 



Steven Tolkin

November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP DOLAP '99**

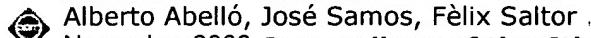
Publisher: ACM Press

Full text available:  pdf(1.23 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the Phoenix system, which loads a data warehouse and then reports against it. Between the raw atomic data of the source system and the business measures presented to users there are many computing environments. Aggregation occurs everywhere: initial bucketing by the natural keys on the mainframe, loading the fact table using a mapping table, maintaining aggregate tables and reporting tables in the data base, in the GUI, in SQL queries issued on behalf of client tools by ...

Keywords: OLAP, SQL, aggregation, data lineage, data warehouse

10 Query processing: Implementing operations to navigate semantic star schemas 



November 2003 **Proceedings of the 6th ACM international workshop on Data warehousing and OLAP DOLAP '03**

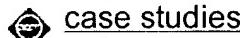
Publisher: ACM Press

Full text available:  pdf(193.82 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In the last years, lots of work have been devoted to multidimensional modeling, star shape schemas and OLAP operations. However, "drill-across" has not captured as much attention as other operations. This operation allows to change the subject of analysis keeping the same analysis space we were using to analyze another subject. It is assumed that this can be done if both subjects share exactly the same analysis dimensions. In this paper, besides the implementation of an algebraic set of operatio ...

Keywords: OLAP operations, SQL, drill-across, semantic relationships, star schema

11 Approaches to the development of multi-dimensional databases: lessons from four 



Helen Hasan, Peter Hyland, David Dodds, Raja Veeraraghavan

May 2000 **ACM SIGMIS Database**, Volume 31 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.21 MB) Additional Information: [full citation](#), [index terms](#)

Keywords: OLAP, case study, executive information systems, management, methodology, multi-multidimensional database, organization, relational database

12 Approaches to the development of multi-dimensional databases: lessons from four case studies

 Helen Hasan, Peter Hyland, David Dodds, Raja Veeraraghavan
June 2000 **ACM SIGMIS Database**, Volume 31 Issue 3

Publisher: ACM Press

Full text available:  pdf(1.21 MB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

The paper explores the manner in which an organization's data and information can be effectively utilized to assist an organization to achieve its business objectives. With the increased popularity of data warehousing and executive information systems, there is renewed interest by IT practitioners in data models and database structures, in particular multi-dimensional forms, which have joined their relational counterparts as legitimate tools for extracting vital business information from an orga ...

Keywords: OLAP, case study, executive information systems, management, methodology, multi-dimensional database, organization, relational database

13 Deriving initial data warehouse structures from the conceptual data models of the underlying operational information systems

 Michael Boehnlein, Achim Ulbrich-vom Ende
November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP DOLAP '99**

Publisher: ACM Press

Full text available:  pdf(1.40 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years the construction of large scale data schemes for operational systems has been the major problem of conceptual data modeling for business needs. Multidimensional data structures used for decision support applications in data warehouses have rather different requirements to data modeling techniques. In case of operational systems the data models are created from application specific requirements. The data models in data warehouses base on the analytical requirements of the use ...

Keywords: conceptual data model, data warehouse, decision support system, entity relationship model (ERM), snowflake scheme, star scheme, structured entity relationship model (SERM)

14 Conceptual modeling for ETL processes

 Panos Vassiliadis, Alkis Simitsis, Spiros Skiadopoulos
November 2002 **Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP DOLAP '02**

Publisher: ACM Press

Full text available:  pdf(471.56 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Extraction-Transformation-Loading (ETL) tools are pieces of software responsible for the extraction of data from several sources, their cleansing, customization and insertion into a data warehouse. In this paper, we focus on the problem of the definition of ETL activities and provide formal foundations for their conceptual representation. The proposed conceptual model is (a) customized for the tracing of inter-attribute relationships and the respective ETL activities in the early stages of a dat ...

Keywords: ETL, conceptual modeling, data warehousing

15 Data Warehouse: Index filtering and view materialization in ROLAP environment

◆ Shi Guang Qiu, Tok Wang Ling

October 2001 **Proceedings of the tenth international conference on Information and knowledge management CIKM '01**

Publisher: ACM Press

Full text available: [pdf\(1.17 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Using materialized view to accelerate OLAP queries is one of the most common methods used in ROLAP systems. However, high storage and computation cost make this method very difficult to be implemented in the actual environment. Among various issues associated with this, index selection and view materialization are two of the top challenges. In this paper, we propose to build indexes on subsets of the primary keys rather than the full sets if the index selectivity for these smaller indexes can be ...

16 An alternative storage organization for ROLAP aggregate views based on cubetrees

◆ Yannis Kotidis, Nick Roussopoulos

June 1998 **ACM SIGMOD Record , Proceedings of the 1998 ACM SIGMOD international conference on Management of data SIGMOD '98**, Volume 27 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.19 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The Relational On-Line Analytical Processing (ROLAP) is emerging as the dominant approach in data warehousing with decision support applications. In order to enhance query performance, the ROLAP approach relies on selecting and materializing in summary tables appropriate subsets of aggregate views which are then engaged in speeding up OLAP queries. However, a straight forward relational storage implementation of materialized ROLAP views is immensely wasteful on storage and incredibly inadeq ...

17 Research session 6: OLAP and constraints: OLAP dimension constraints

◆ Carlos A. Hurtado, Alberto O. Mendelzon

June 2002 **Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems PODS '02**

Publisher: ACM Press

Full text available: [pdf\(186.06 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In multidimensional data models intended for online analytic processing (OLAP), data are viewed as points in a multidimensional space. Each dimension has structure, described by a directed graph of categories, a set of members for each category, and a child/parent relation between members. An important application of this structure is to use it to infer summarizability, that is, whether an aggregate view defined for some category can be correctly derived from a set of precomputed views defined f ...

18 Data warehouse construction: Building a web warehouse for accessibility data

◆ Christian Thomsen, Torben Bach Pedersen

November 2006 **Proceedings of the 9th ACM international workshop on Data warehousing and OLAP DOLAP '06**

Publisher: ACM Press

Full text available: [pdf\(181.96 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

As more and more information is available on the web, it is a problem that many web resources are not accessible, i.e., are not usable for users with special needs. For

example, for a web page to be accessible, it should give text alternatives (i.e., explanatory texts) for images such that blind users that have the web pages read aloud automatically also can obtain information about the images. In the European Internet Accessibility Observatory (EIAO) project, a crawler that will evaluate the ac ...

Keywords: accessibility, web warehouse

19 Reports: Report on the ACM fourth international workshop on data warehousing and 

 **OLAP (DOLAP 2001)**

Joachim Hammer

June 2002 **ACM SIGMOD Record**, Volume 31 Issue 2

Publisher: ACM Press

Full text available:  pdf(413.34 KB) Additional Information: [full citation](#), [abstract](#)

The Fourth Annual ACM International Workshop on Data Warehousing and Online Analytical Processing (DOLAP 2001) was held in Atlanta, GA, USA, in November 2001, in conjunction with the Tenth International Conference on Information and Knowledge Management (CIKM 2001). Although this was only the fourth annual meeting, DOLAP has already become an important and broadly accepted forum for researchers and practitioners to share their findings in theoretical foundations, current methodologies, practical ...

20 A method for developing dimensional data marts 

 Tim Chenoweth, David Schuff, Robert St. Louis

December 2003 **Communications of the ACM**, Volume 46 Issue 12

Publisher: ACM Press

Full text available:  pdf(104.91 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index](#)
 html(26.61 KB) [terms](#)

Decision-oriented dimensional data marts are fundamentally different than transaction-oriented relational databases. A distinctive methodology and a different set of tools are required for their effective development.

Results 1 - 20 of 25

Result page: **1** [2](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

091954,719


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
 The ACM Digital Library The Guide

 "selecting web" + "dimension table" + "specifying query" + "fa

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [selecting web](#) [dimension table](#) [specifying query](#) [fact table](#) [produce join result](#)

Found 168 of 198,617

Sort results by Save results to a Binder
 Search Tips
 Display results Open results in a new window

[Try an Advanced Search](#)
[Try this search in The ACM Guide](#)

Results 1 - 20 of 168

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [next](#)

Relevance scale

- 1 [Query processing: Exploiting hierarchical clustering in evaluating multidimensional aggregation queries](#)

Dimitri Theodoratos
 November 2003 **Proceedings of the 6th ACM international workshop on Data warehousing and OLAP DOLAP '03**

Publisher: ACM Press

Full text available: [pdf\(216.79 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Multidimensional aggregation queries constitute the single most important class of queries for data warehousing applications and decision support systems. The bottleneck in the evaluation of these queries is the join of the usually huge fact table with the restricted dimension tables (*star-join*). Recently, a multidimensional hierarchical clustering schema for star schemas is suggested. Subsequently, query evaluation plans for multidimensional queries appeared that essentially implement a ...

Keywords: multidimensional aggregation query, multidimensional hierarchical clustering, query transformations, star join

- 2 [Heuristic optimization of OLAP queries in multidimensionally hierarchically clustered databases](#)

Dimitri Theodoratos, Aris Tsois
 November 2001 **Proceedings of the 4th ACM international workshop on Data warehousing and OLAP DOLAP '01**

Publisher: ACM Press

Full text available: [pdf\(1.44 MB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

On-line analytical processing (OLAP) is a technology that encompasses applications requiring a multidimensional and hierarchical view of data. OLAP applications often require fast response time to complex grouping/aggregation queries on enormous quantities of data. Commercial relational database management systems use mainly multiple one-dimensional indexes to process OLAP queries that restrict multiple dimensions. However, in many cases, multidimensional access methods outperform one-dimensional ...

- 3 [On the problem of generating common predecessors](#)

W. Lehner, W. Hümmer, L. Schlesinger, A. Bauer

November 2002 **Proceedings of the 5th ACM international workshop on Data Warehousing and OLAP DOLAP '02**

Publisher: ACM Press

Full text available: [pdf\(231.16 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Using common subexpressions to speed up a set of queries is a well known and long studied problem. However, due to the isolation requirement, operating a database in the classic transactional way does not offer many applications to exploit the benefits of simultaneously computing a set of queries. In the opposite, many applications can be identified in the context of data warehousing, e. g. optimizing the incremental maintenance process of multiple dependent materialized views or the generation ...

4 **Maintenance of data cubes and summary tables in a warehouse**

Inderpal Singh Mumick, Dallan Quass, Barinderpal Singh Mumick

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available: [pdf\(1.58 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Data warehouses contain large amounts of information, often collected from a variety of independent sources. Decision-support functions in a warehouse, such as on-line analytical processing (OLAP), involve hundreds of complex aggregate queries over large volumes of data. It is not feasible to compute these queries by scanning the data sets each time. Warehouse applications therefore build a large number of summary tables, or materialized aggregate views, to ...

5 **Physical design: Handling big dimensions in distributed data warehouses using the DWS technique**

Marco Costa, Henrique Madeira

November 2004 **Proceedings of the 7th ACM international workshop on Data warehousing and OLAP DOLAP '04**

Publisher: ACM Press

Full text available: [pdf\(288.33 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The DWS (Data Warehouse Striping) technique allows the distribution of large data warehouses through a cluster of computers. The data partitioning approach partitions the facts tables through all nodes and replicates the dimension tables. The replication of the dimension tables creates a limitation to the applicability of the DWS technique to data warehouses with big dimensions. This paper proposes a strategy to handle large dimensions in a distributed DWS system and evaluates the proposed str ...

Keywords: data warehousing, distributed query execution

6 **Automated data warehousing for rule-based CRM systems**

Han-joon Kim, TaeHee Lee, Sang-goo Lee, Jonghun Chun

January 2003 **Proceedings of the 14th Australasian database conference - Volume 17 ADC '03**

Publisher: Australian Computer Society, Inc.

Full text available: [pdf\(274.28 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper proposes a novel way of automatically developing data warehouse configuration in rule-based CRM systems. Rule-based CRM systems assume that marketing activities are represented as a set of **IF-WHEN** rules. Currently, to provide good quality CRM functionalities, CRM systems seek to combine conventional CRM methodologies with data warehousing technology. A data warehouse can be abstractly seen as a set of materialized views. Selecting views for materialization in a data

warehouse i ...

Keywords: CRM, analysis query, data warehouse, materialized view, rules, star-join index

7 Industrial sessions: commercial implementation techniques: Efficient execution of

 joins in a star schema

Andreas Weininger

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available:  pdf(349.43 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

A star schema is very popular for modeling data warehouses and data marts. Therefore, it is important that a database system which is used for implementing such a data warehouse or data mart is able to efficiently handle operations on such a schema. In this paper we will describe how one of these operations, the join operation --- probably the most important operation --- is implemented in the IBM Informix Extended Parallel Server (XPS).

8 Industrial sessions: big data: TPC-DS, taking decision support benchmarking to the

 next level

Meikel Poess, Bryan Smith, Lubor Kollar, Paul Larson

June 2002 **Proceedings of the 2002 ACM SIGMOD international conference on Management of data SIGMOD '02**

Publisher: ACM Press

Full text available:  pdf(645.38 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

TPC-DS is a new decision support benchmark currently under development by the Transaction Processing Performance Council (TPC). This paper provides a brief overview of the new benchmark. The benchmark models the decision support functions of a retail product supplier, including data loading, multiple types of queries and data maintenance. The database consists of multiple snowflake schemas with shared dimension tables; data is skewed; and the query set is large. Overall, the benchmark is consider ...

Keywords: TPC, benchmark, data warehouse, decision support, performance evaluation

9 Dynamic maintenance of multidimensional range data partitioning for parallel data

 processing

Junping Sun, William I. Grosky

November 1998 **Proceedings of the 1st ACM international workshop on Data warehousing and OLAP DOLAP '98**

Publisher: ACM Press

Full text available:  pdf(1.09 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

10 Dealing with slow-evolving fact: a case study on inventory data warehousing

 Chung-Min Chen, Munir Cochinwala, Elsa Yueh

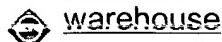
November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP DOLAP '99**

Publisher: ACM Press

Full text available:  pdf(941.25 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Data Warehousing for INventory management (DWIN) is a production project at Telcordia aimed at providing telecommunications service providers with decision support functions for inventory control and monitoring. In this paper, we report some interesting issues related to the design of the data warehouse. Specifically, we will discuss the issues of slow-evolving fact, transaction-oriented fact table, and large dimensions. We also propose the concept of virtual data cubes and ...

11 Aggregation everywhere: data reduction and transformation in the Phoenix data



Steven Tolkin

November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP DOLAP '99**

Publisher: ACM Press

Full text available: [pdf\(1.23 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes the Phoenix system, which loads a data warehouse and then reports against it. Between the raw atomic data of the source system and the business measures presented to users there are many computing environments. Aggregation occurs everywhere: initial bucketing by the natural keys on the mainframe, loading the fact table using a mapping table, maintaining aggregate tables and reporting tables in the data base, in the GUI, in SQL queries issued on behalf of client tools by ...

Keywords: OLAP, SQL, aggregation, data lineage, data warehouse

12 Optimizing large star-schema queries with snowflakes via heuristic-based query rewriting

Yingying Tao, Qiang Zhu, Calisto Zuzarte, Wing Lau

October 2003 **Proceedings of the 2003 conference of the Centre for Advanced Studies on Collaborative research CASCON '03**

Publisher: IBM Press

Full text available: [pdf\(182.43 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

User queries have been becoming increasingly complex (e.g., involving a large number of joins) as database technology is applied to some application domains such as data warehouses and life sciences. Query optimizers in existing database management systems often suffer from intolerably long optimization time and/or poor optimization results when optimizing large join queries. One possible solution to tackle these problems is to rewrite a user-specified complex query into another form that can be ...

Keywords: complex query, database management system, query graph, query optimization, query rewrite

13 Query processing: Implementing operations to navigate semantic star schemas

Alberto Abelló, José Samos, Fèlix Saltor

November 2003 **Proceedings of the 6th ACM international workshop on Data warehousing and OLAP DOLAP '03**

Publisher: ACM Press

Full text available: [pdf\(193.82 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In the last years, lots of work have been devoted to multidimensional modeling, star shape schemas and OLAP operations. However, "drill-across" has not captured as much attention as other operations. This operation allows to change the subject of analysis keeping the same analysis space we were using to analyze another subject. It is assumed that this can be done if both subjects share exactly the same analysis dimensions. In this

paper, besides the implementation of an algebraic set of operations ...

Keywords: OLAP operations, SQL, drill-across, semantic relationships, star schema

14 Optimizing multiple dimensional queries simultaneously in multidimensional databases

Weifa Liang, Maria E. Orlowska, Jeffrey X. Yu

February 2000 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 8 Issue 3-4

Publisher: Springer-Verlag New York, Inc.

Full text available:  pdf(269.57 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Some significant progress related to multidimensional data analysis has been achieved in the past few years, including the design of fast algorithms for computing data cubes, selecting some precomputed group-bys to materialize, and designing efficient storage structures for multidimensional data. However, little work has been carried out on multidimensional query optimization issues. Particularly the response time (or evaluation cost) for answering several related dimensional queries simultaneous ...

Keywords: Data warehousing, MDDBs, Multiple dimensional query optimization, OLAP, Query modeling

15 Why decision support fails and how to fix it

 Ralph Kimball, Kevin Strehlo

September 1995 **ACM SIGMOD Record**, Volume 24 Issue 3

Publisher: ACM Press

Full text available:  pdf(722.24 KB) Additional Information: [full citation](#), [citations](#), [index terms](#)

16 Improved query performance with variant indexes

 Patrick O'Neil, Dallan Quass

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data SIGMOD '97**, Volume 26 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.54 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The read-mostly environment of data warehousing makes it possible to use more complex indexes to speed up queries than in situations where concurrent updates are present. The current paper presents a short review of current indexing technology, including row-set representation by Bitmaps, and then introduces two approaches we call Bit-Sliced indexing and Projection indexing. A Projection index materializes all values of a column in RID order, and a Bit-Sliced index essentially takes an orth ...

17 An overview of data warehousing and OLAP technology

 Surajit Chaudhuri, Umeshwar Dayal

March 1997 **ACM SIGMOD Record**, Volume 26 Issue 1

Publisher: ACM Press

Full text available:  pdf(101.60 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

Data warehousing and on-line analytical processing (OLAP) are essential elements of decision support, which has increasingly become a focus of the database industry. Many commercial products and services are now available, and all of the principal database management system vendors now have offerings in these areas. Decision support places

some rather different requirements on database technology compared to traditional on-line transaction processing applications. This paper provides an overview ...

18 RSS and views: Fast approximate computation of statistics on views

 Calisto Zuzarte, Xiaohui Yu
June 2006 **Proceedings of the 2006 ACM SIGMOD international conference on Management of data SIGMOD '06**

Publisher: ACM Press

Full text available: [pdf\(38.25 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Accurate estimation of the sizes of intermediate query results (cardinality estimation) is of critical importance to plan costing in query optimization. The common practice in current commercial database systems such as IBM DB2 Universal Database (DB2 UDB) is to derive the cardinality estimates from base-table statistics. However, this approach often suffers from simplifying yet unrealistic assumptions that have to be made about the underlying data (for example, different attributes are independ ...

Keywords: query optimization, selectivity estimation

19 Deriving initial data warehouse structures from the conceptual data models of the

 underlying operational information systems

Michael Boehnlein, Achim Ulbrich-vom Ende

November 1999 **Proceedings of the 2nd ACM international workshop on Data warehousing and OLAP DOLAP '99**

Publisher: ACM Press

Full text available: [pdf\(1.40 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years the construction of large scale data schemes for operational systems has been the major problem of conceptual data modeling for business needs. Multidimensional data structures used for decision support applications in data warehouses have rather different requirements to data modeling techniques. In case of operational systems the data models are created from application specific requirements. The data models in data warehouses base on the analytical requirements of the use ...

Keywords: conceptual data model, data warehouse, decision support system, entity relationship model (ER), snowflake scheme, star scheme, structured entity relationship model (SERM)

20 XPS a database server for data warehousing

 Andreas Weininger
November 2001 **Proceedings of the 4th ACM international workshop on Data warehousing and OLAP DOLAP '01**

Publisher: ACM Press

Full text available: [pdf\(621.63 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

A database server used for implementing a data warehouse must support other features than a database server used for OLTP. Therefore, in this paper we will look specifically at features necessary for efficiently processing queries on a database with a star schema model, a database scheme which is used very often in data warehousing. We will especially analyze the features provided for this by the IBM Extended Parallel Server (XPS). There are special star join methods like the Push-Down Hash Semi ...

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)

09/954,719


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 Search Results**BROWSE****SEARCH****IEEE XPLORE GUIDE**

Results for "(((dimension* <near> table*) <paragraph> (fact <near> table*) <paragraph> qu...)"
 Your search matched 9 of 1521373 documents.

e-mail

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.

» Search Options[View Session History](#)[New Search](#)**Modify Search**

[Search](#)
 Check to search only within this results set

Display Format: [Citation](#) [Citation & Abstract](#)
[view selected items](#) [Select All](#) [Deselect All](#)
» Key
IEEE JNL IEEE Journal or Magazine

IET JNL IET Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IET CNF IET Conference Proceeding

IEEE STD IEEE Standard
 1. Web-enabled Temporal OLAP

Vaisman, A.A.; Izquierdo, A.; Ktenas, M.;
[Web Congress, 2006. LA-Web '06. Fourth Latin American Oct. 2006](#) Page(s):220 - 229
 Digital Object Identifier 10.1109/LA-WEB.2006.37

[AbstractPlus](#) | Full Text: [PDF\(259 KB\)](#) [IEEE CNF Rights and Permissions](#)

 2. Experiments on Remote Sensing image cube and its OLAP

Mingjie Xu;
[Geoscience and Remote Sensing Symposium, 2004. IGARSS '04. Proceeding International](#)
 Volume 7, 2004 Page(s):4398 - 4401 vol.7
 Digital Object Identifier 10.1109/IGARSS.2004.1370124

[AbstractPlus](#) | Full Text: [PDF\(349 KB\)](#) [IEEE CNF Rights and Permissions](#)

 3. Ad-hoc association-rule mining within the data warehouse

Nestorov, S.; Jukic, N.;
[System Sciences, 2003. Proceedings of the 36th Annual Hawaii International C](#)
 6-9 Jan 2003 Page(s):10 pp.
 Digital Object Identifier 10.1109/HICSS.2003.1174605

[AbstractPlus](#) | Full Text: [PDF\(439 KB\)](#) [IEEE CNF Rights and Permissions](#)

 4. Combining hierarchy encoding and pre-grouping: intelligent grouping in processing

Pieringer, R.; Elhardt, K.; Ramsak, F.; Markl, V.; Fenk, R.; Bayer, R.; Karayannellis, T.;
[Data Engineering, 2003. Proceedings. 19th International Conference on](#)
 5-8 March 2003 Page(s):329 - 340

[AbstractPlus](#) | Full Text: [PDF\(711 KB\)](#) [IEEE CNF Rights and Permissions](#)

 5. An efficient processing of queries with joins and aggregate functions in a warehousing environment

Jin-Ho Kim; Yun-Ho Kim; Sang-Wook Kim; Soo-Ho Ok;
[Database and Expert Systems Applications, 2002. Proceedings. 13th Internati](#)

2-6 Sept. 2002 Page(s):785 - 791

[AbstractPlus](#) | Full Text: [PDF\(500 KB\)](#) IEEE CNF
[Rights and Permissions](#)

6. **Exploitation of pre-sortedness for sorting in query processing: the Temp' for UB-trees**
Zirkel, M.; Markl, V.; Bayer, R.;
[Database Engineering & Applications, 2001 International Symposium on](#),
16-18 July 2001 Page(s):155 - 166
Digital Object Identifier 10.1109/IDEAS.2001.938082
[AbstractPlus](#) | Full Text: [PDF\(1044 KB\)](#) IEEE CNF
[Rights and Permissions](#)
7. **What can partitioning do for your data warehouses and data marts?**
Bellatreche, L.; Karlapalem, K.; Mohania, M.; Schneider, M.;
[Database Engineering and Applications Symposium, 2000 International](#),
18-20 Sept. 2000 Page(s):437 - 445
Digital Object Identifier 10.1109/IDEAS.2000.880634
[AbstractPlus](#) | Full Text: [PDF\(720 KB\)](#) IEEE CNF
[Rights and Permissions](#)
8. **Management of multidimensional aggregates for efficient online analytics**
Albrecht, J.; Bauer, A.; Deyerling, O.; Gunzel, H.; Hummer, W.; Schlesinger, H.;
[Database Engineering and Applications, 1999. IDEAS '99. International Symposium on Proceedings](#),
2-4 Aug. 1999 Page(s):156 - 164
Digital Object Identifier 10.1109/IDEAS.1999.787264
[AbstractPlus](#) | Full Text: [PDF\(144 KB\)](#) IEEE CNF
[Rights and Permissions](#)
9. **Data warehouse design for pharmaceutical drug discovery research**
Axel, M.G.; Il-Yeol Song;
[Database and Expert Systems Applications, 1997. Proceedings., Eighth International Conference on](#),
1-2 Sept. 1997 Page(s):644 - 650
Digital Object Identifier 10.1109/DEXA.1997.617399
[AbstractPlus](#) | Full Text: [PDF\(564 KB\)](#) IEEE CNF
[Rights and Permissions](#)


[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

 Search Results**BROWSE****SEARCH****IEEE XPLOR GUIDE**

Results for "(((dimension* <near> table*) <paragraph> (fact <near> table*) <and> (query* ...")
 Your search matched 3 of 1521373 documents.

 [e-mail](#)

A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.

» **Search Options**[View Session History](#)[New Search](#)**Modify Search**

 [Search](#) Check to search only within this results setDisplay Format: Citation Citation & Abstract» **Key**

IEEE JNL IEEE Journal or Magazine

IET JNL IET Journal or Magazine

IEEE CNF IEEE Conference Proceeding

IET CNF IET Conference Proceeding

IEEE STD IEEE Standard

 [view selected items](#)[Select All](#) [Deselect All](#)

- 1. **Experiments on Remote Sensing image cube and its OLAP**
 Mingjie Xu;
Geoscience and Remote Sensing Symposium, 2004. IGARSS '04. Proceeding International
 Volume 7, 2004 Page(s):4398 - 4401 vol.7
 Digital Object Identifier 10.1109/IGARSS.2004.1370124
AbstractPlus | Full Text: [PDF\(349 KB\)](#) [IEEE CNF](#)
Rights and Permissions

- 2. **Combining hierarchy encoding and pre-grouping: intelligent grouping in processing**
 Pieringer, R.; Elhardt, K.; Ramsak, F.; Markl, V.; Fenk, R.; Bayer, R.; Karayann Sellis, T.;
Data Engineering, 2003. Proceedings. 19th International Conference on
 5-8 March 2003 Page(s):329 - 340
AbstractPlus | Full Text: [PDF\(711 KB\)](#) [IEEE CNF](#)
Rights and Permissions

- 3. **Exploitation of pre-sortedness for sorting in query processing: the Temp for UB-trees**
 Zirkel, M.; Markl, V.; Bayer, R.;
Database Engineering & Applications, 2001 International Symposium on
 16-18 July 2001 Page(s):155 - 166
 Digital Object Identifier 10.1109/IDEAS.2001.938082
AbstractPlus | Full Text: [PDF\(1044 KB\)](#) [IEEE CNF](#)
Rights and Permissions

[Help](#) [Contact Us](#) [Privacy &](#)

© Copyright 2006 IEEE -

 Indexed by